

Automatic 3D underground mine mapping

Daniel F. Huber Nicolas Vandapel

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213
dhuber@ri.cmu.edu vandapel@ri.cmu.edu

Abstract

For several years, our research group has been developing methods for automated modeling of 3D environments. In September, 2002, we were given the opportunity to demonstrate our mapping capability in an underground coal mine. The opportunity arose as a result of the Quecreek mine accident, in which an inaccurate map caused miners to breach an abandoned, water-filled mine, trapping them for several days. Our field test illustrates the feasibility and potential of high resolution three-dimensional (3D) mapping of an underground coal mine using a cart-mounted 3D laser scanner. This paper presents our experimental setup, the automatic 3D modeling method used, and the results of the field test. In addition, we address issues related to laser sensing in a coal mine environment.

1 Introduction

For several years, our research group has been developing methods for automated modeling of 3D environments [3][4][5]. In September, 2002, we were given the opportunity to demonstrate our mapping capability in an underground coal mine, the Mine Safety and Health Administration (MSHA) research mine in Bruceton, Pennsylvania. The opportunity arose as a result of the Quecreek mine accident in July, 2002, in which miners inadvertently breached an abandoned, water-filled mine, trapping themselves amidst thousands of tons of water. After the miners were safely rescued, an investigation was launched to determine the cause of the accident and to identify new procedures necessary to prevent mine breaches in the future. Regulations already in place aim to prevent such an accident: mapping the mine before ending operations, exploratory drilling, and so forth. Unfortunately, old maps may be incorrect, incomplete, or simply lost. In the end, the Quecreek accident was attributed to an inaccurate map [2].

A collaborative effort by several research groups at Carnegie Mellon University (CMU) has been formed to develop robots to autonomously map abandoned mines and active mines before operations are ended. Such robots would be an important contribution to mining safety. Details can be found in [1, 14, 9]. In this paper, we address the problem of sensing and generating high-resolution 3D models of an active mine. In September, 2002, we conducted a field test to demonstrate the feasibility of high resolution 3D mapping of an underground coal mine using a cart-mounted 3D laser scanner. The remainder of the paper is organized as follows. First, section 2 reviews previous work on mine mapping and localization. Section 3 describes our experimental setup and the data collection process. Section 4 explains our automatic modeling algorithm and discusses the resulting 3D model. Finally, section 5 presents an analysis of the issues relevant to laser sensing in a coal mine environment.

2 Related work

In this section, we review the most relevant work on mine mapping and localization. Early work by Shaffer [13] described a method to localize a mobile robot in an underground mine by registering terrain features (corner and line segments) extracted from an *a priori* survey map with cross-sections from an environment map produced by a laser scanner. In [11, 12], Scheduling extensively tested a set of navigation sensors mounted on a Load, Haul, and Dump truck (LHD) in the harsh underground mine environment. Using the data from a laser line scanner coupled with the navigation data of the vehicle, he produced a 3D model of a section of the mine. In [7], two line scanners were integrated on an LHD. The iterative closest point (ICP) algorithm was used to register the 2D profiles to an existing map. This implementation was extended to mine mapping in [8]. The contributions presented above focused on vehicle automation for active mines. In the context of mapping



Figure 1: The cart-mounted Z+F laser scanner used in the data collection.

abandoned mines, Thrun [14] produced 2D maps and partial 3D models of tunnels, using a SLAM approach with two line scanning lasers mounted on a tele-operated robot [1]. Several systems have been designed to map mines that are inaccessible to a ground robot, for example, by mapping a cavity using a 3D laser sensor inserted through a bore-hole. Such systems include the C-ALS (Cavity Autoscanning laser system) by Measurement Devices, Ltd. and the Cavity monitoring system by Optech, Inc.¹ A similar approach has been followed in [9].

3 Data collection

For our field test, we used a high resolution 3D laser scanner mounted on a cart as illustrated in figure 1. The sensor, a Zoller and Fröhlich LARA 25200 (Z+F) scanner [6], produces 8000×1400 pixel range and reflectance images with millimeter-level accuracy. The field of view is $360^\circ \times 70^\circ$ with a range of 22.5 m. The laser scan head was inclined to allow higher density scanning of the floor and ceiling near the scanning platform. Unfortunately, in some regions, the low roof was actually too close to the scan head for the sensor to fully scan the ceiling.

We obtained 23 scans at three- to five-meter intervals along a loop trajectory through a sequence of 4 hallways (figure 2). The cart was kept stationary at each location for the 90 seconds required to obtain each scan. Due to the capabilities of our modeling algorithms, it was not necessary to record the position or attitude of the cart. This greatly simplifies the data collection process. The entire procedure only took about three hours, including setup and disassembly of the equipment. For this experiment, the cart was moved manually, but it would be straightforward to mount the scanner on an autonomous mobile robot.

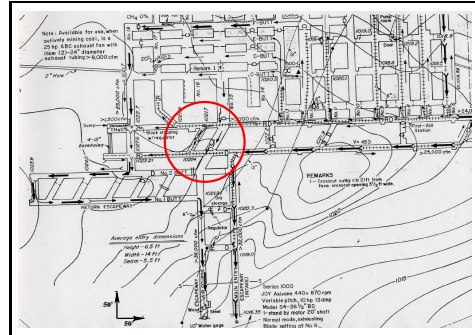


Figure 2: Surveyed map of the Bruceton mine. The red circle indicates the area mapped in our field experiment.

4 Automatic modeling from reality

Modeling-from-reality is the process of creating digital three-dimensional (3D) models of real-world scenes from 3D views as obtained, for example, from range sensors or stereo camera systems. Recently, we have developed a system that fully automates the modeling-from-reality process [4][5]. The key challenge of automatic modeling-from-reality is the accurate and robust registration of multiple 3D views. Although each input scan is an accurate representation of the 3D structure of the scene as seen from a single viewpoint, the data is expressed in the local coordinate system of the sensor. Our system automatically registers multiple 3D data sets in a common coordinate system without requiring any knowledge of the viewpoints from which the data was obtained. This capability is important in our case, because we did not survey the scan locations during our initial data collection. In a real system, where the sensor would be mounted on a robot, an approximate estimate of the motion between scans may be provided by the robot. Our algorithm has the ability to employ such information when it is available, but, more importantly, it will not break down when the estimates are not available. The problem of registering multiple 3D views obtained from unknown viewpoints is called multi-view surface matching, and it is analogous to assembling a 3D jigsaw puzzle, with each view being a piece of the puzzle.

4.1 Automatic modeling method

Briefly, our automatic modeling-from-reality algorithm works as follows²: First, the range images are converted to 3D surfaces, which serve as the input views to the algorithm. A pair-wise surface matching algorithm is then used to identify potential alignments (called matches) between view pairs. The difficulty is that these pair-wise matches

¹www.mdl.co.uk, www.optech.on.ca, May 2003

²Details of the algorithm are given in [4].




Figure 3: Coal mine model graphs. Matches are hand-labeled for illustration: thick (blue) edges are correct matches, and thin (red) edges are incorrect matches. (a) The model graph from exhaustive pair-wise registration; (b) After filtering the matches for local consistency (G_{LR}). (c) The solution found by our multi-view surface matching algorithm. (d) The solution with all overlapping views used for verification. Here, the position of the nodes represents the x,y position of each view as determined by our algorithm.

may be incorrect, and it is sometimes impossible to distinguish correct matches from incorrect ones just looking at pairs of views. We overcome this problem by evaluating the quality of an entire network of matches. This network is a graph data structure called the model graph (figure 3). It contains a node for each input view and an edge for each pair-wise match.

Figure 3-(a) shows the model graph for exhaustive pair-wise registration of the views from the mine data set. Matches are hand-labeled as correct or incorrect for illustration³. The worst pair-wise matches are then removed using a local consistency test (figure 3-(b)). We call this model graph G_{LR} (LR stands for local registration). If we can find a connected sub-graph of G_{LR} that contains only correct matches, it is straightforward to convert the relative poses associated with the matches into absolute poses for each view. We formulate this search as a mixed continuous and discrete optimization problem. The discrete optimization performs a combinatorial search over the space of connected sub-graphs of the model graph, using a global consistency measure to detect and avoid incorrect, but locally consistent matches. The continuous optimization adjusts the absolute pose parameters to minimize the registration error between all overlapping surfaces, distributing the pair-wise registration errors in a principled way.

Unfortunately, a connected sub-graph containing only correct matches may not exist within G_{LR} . This could happen if one of the input views is corrupted or if some subset of the input views does not overlap sufficiently with any of the remaining views (e.g., if the sensor is moved too far between two scans). We have developed a class of search algorithms that handles these situations gracefully by searching the space of *all* sub-graphs of G_{LR} for the *best* model

hypothesis rather than searching only for globally consistent connected sub-graphs.

The output of the optimization is a set of rigid body transforms that aligns the views in a common coordinate system. If the best model hypothesis is not a connected sub-graph of G_{LR} , each component of the sub-graph is a separate part of a multi-part model, and the algorithm outputs transforms that register the views within each part. The redundant surfaces in the registered views can then be merged to form a more compact, unified 3D model.

We have developed several different search algorithms for finding the best model hypothesis within G_{LR} . The method used in this paper is called iterative merging. This algorithm begins with a model hypothesis in which every view is a separate part (i.e., a model graph with a node for each view and no edges). Each iteration of the al-



Figure 4: Perspective view of the automatically constructed mine model. The model consists of 23 separate views covering a region 43x42 m wide and 3 m high.

³A match is defined to be correct if no point in either scan is displaced more than 10% of the scene size from its correct location.



Figure 5: Close-up views of the 3D model from various viewpoints (top row) and photographs obtained from the same viewpoints (bottom row).

gorithm merges two parts using the relative pose from a pair-wise match. The chances of merging two parts using an incorrect match are minimized by merging the best-matching parts first and by verifying that all of the views in the merged part are mutually consistent. The model graph produced by the iterative merging algorithm is shown in figure 3-(c) and (d). The corresponding 3D model is shown in figure 4.

4.2 Discussion

Figure 5 shows three close-up views of the 3D model along with photographs taken from the same locations. The details visible in the model highlight the quality possible with high-resolution 3D scanning. In the left picture, a pair of doors with z-braces can be seen. In the center picture, stacks of bricks lie against the walls and on pallets. The partially buried track rails run the length of the corridor and a junction in the rails can be seen in the distance. In the right picture, the rails are more pronounced, and several large pipes are stored on both sides of the corridor.

The difference between the surveyed map of the mine and the actual mine can be significant, as shown in figure 6. The surveyed map is topologically correct, but the scale in parts of the map is significantly off. For example, the trapezoidal column of rock in the center of the surveyed map is actually much narrower near the bottom of the map, which might be a safety concern. The 3D model allows direct measurement of the wall thickness at any point. Furthermore, the 3D model enables measurements that are not possible on the 2D map. For example, it is straightforward to compute the ceiling height throughout the mine or to accurately measure the amount of material that has been removed from the mine.

Although we do not have ground truth measurements of the sensor positions, we can still perform some analysis of the accuracy of the model. First, we can tell that the sensor positions are qualitatively correct by looking at cross sections of the model (e.g., figure 6-(b)). If any view were incorrectly registered, some surfaces would not align well with one another and would appear as errant surfaces in the cross sections. The distance between overlapping surfaces gives a quantitative measure of the model accuracy. For this data, the RMS distance between overlapping surfaces was 1.09 cm. Finally, we have performed analysis of other data sets of environments of a similar size for which we measured ground truth positions using a theodolite and fiducials ([4], p. 160). In those experiments, we found the pose error to be less than 0.25% of the model size. The RMS error of the fiducial marker measurements using our algorithm was comparable to the error in the same measurements using bundle-block adjustment on the fiducial points (2.7 cm RMS versus 2.5 cm RMS). This suggests that the primary source of error in that experiment was the




Figure 6: (a) A close-up view of the map from figure 2 showing the region mapped in our field test. (b) A horizontal cross-section of the 3D model of the corresponding area.

measurement of the fiducial position rather than a limitation of our multi-view surface matching algorithm.

5 Laser sensing in mine environments

Coal mine environments present a number of unique challenges for laser sensing systems, including the presence of explosive gas, widely ranging surface albedo, metallic objects, and wet surfaces. In the field test mine, the walls were coated non-uniformly with a white, waterproofing material, and in many places, bare coal was exposed. The roof was reinforced with metallic netting, and the environment contained numerous metallic objects, such as pipes and rails. Furthermore, regions of the walls and ceiling were wet and dripping water. For additional experiments, we collected samples of rocks and bituminous coal for analysis in the controlled environment of our laboratory.

Open beam lasers can be a potential ignition source of methane gas or coal dust, but studies have shown that below 150 mW or 20 mW/mm² methane gas or coal dust cannot be ignited by a laser beam⁴. With an average power of 22 mW [6], the Z+F laser poses no threat.

Our second concern was the level of noise and bias in range measurements when scanning scenes with widely-ranging surface albedo. We analyzed the noise and bias using a calibration target made of 6 different color patches, including black and white (figure 7-(a)). We positioned the laser at 7.5 m from the target and collected 11 identical scans to test the repeatability of the measurement. We measured the range for the pixels within each patch (725 pixels) and computed the mean and standard deviation for each patch over all the scans. Figure 7-(b) shows the distribution of range measurements for the black and white patches. As expected, the level of noise for the black target ($\sigma = 5.80$ mm) is larger than that of the white target ($\sigma = 3.54$ mm); however, even the worst case noise, which occurred with the black patch, was acceptable for a mine-mapping application. We analyzed reflectance-based range bias by estimating the difference in range between the white and black patch. For this experiment, we fit planes to the two patches using the total least squares method. Figure 7-(c) shows a top view of the two estimated planes, which have an offset of 1.3 cm. As with the noise error, this bias is within acceptable limits for mine-mapping.

Finally, we considered the effect of scanning specular targets, such as bituminous coal (which is relatively shiny) or wet surfaces. To test this, we scanned a flat piece of coal twice – once when the sample was dry and again when wet. The sample was positioned at 7.5 m from the sensor and scanned at near-normal incidence. As expected

| Sample | Reflectance (min-mean-max) | | |
|-------------|----------------------------|------|------|
| Coal (dry) | 152 | 411 | 1347 |
| Coal (wet) | 26 | 145 | 486 |
| Rock | 291 | 491 | 754 |
| Wood | 4413 | 5663 | 6709 |
| Aluminum | 3330 | 3698 | 4022 |
| Black paper | 283 | 355 | 470 |
| White paper | 5353 | 5571 | 5777 |

Table 1: Reflectance for different targets at 7.5 m

the dry sample produced erroneous range measurements associated with specular reflections. Surprisingly, the wetting the coal sample actually reduced the frequency of erroneous measurements. We hypothesize that the reason we did not experience many specular reflections in our field tests is due to the wall-coating and damp environment. Table 1 shows a comparison of reflectance values for several targets scanned at 7.5 m and near-normal incidence, including the wet and dry coal samples.

6 Summary and future work

In this paper, we have shown that our automatic modeling-from-reality algorithms can be successfully applied to the problem of high-resolution mapping of underground mines. The model constructed from the 23 scans obtained during our field test was estimated to contain geometric errors on the order of 1 cm. The results of our laboratory experiments indicate that the various sensing challenges presented by the underground mining scenario may introduce error of 1-2 cm into a 3D model. However, it should be noted that these tests are only partially representative because the environment in our laboratory and in the Bruceton coal mine do not fully mimic harsh environment of an active coal mine.

The results of this paper are a proof of concept. The next step would be to further specialize our automatic modeling system for the purpose of mine mapping. First, a ruggedized platform for the system must be developed, either in the form of an electric cart or a tele-operated mobile robot. Second, our automatic modeling algorithms should be modified to operate in an online mode as opposed to the current batch method. The immediate feedback of an online algorithm would enable mine mappers to effectively plan the scan locations. Finally, we are working on new modeling algorithms that scale to very large numbers of views. Our current algorithms have $O(N^2)$ complexity in the number of input views, which limits processing to sub-maps containing about 50 views.

Acknowledgments

We would like to thank the Pittsburgh Research Laboratory of the National Institute for Occupational Safety and

⁴NIOSH, May 2003, Laser Safety in Underground Mine, www.cdc.gov/niosh/mining/lasersafety/default.htm




Figure 7: (a) Photograph of the calibration target with coal samples in the foreground. (b) Distribution of range measurements for the white and black targets. (c) Top view of the two planes fit to the black and white patches showing the reflectance-based range bias.

Health (NIOSH) for the access to the Bruceon research mine. This research has been supported in part by a fellowship from the Eastman Kodak Company and by the National Science Foundation under grant 0102272.

References

- [1] C. Baker, Z. Omohundro, S. Thayer, W. Whittaker, M. Montemerlo and S. Thrun, "Case Studies in Robotic Mine Mapping", *International Conference on Field and Service Robotics*, 2003.
- [2] T. Gibb and D. Hopey, "Quecreek mine accident report blames outdated map", *Pittsburgh Post-Gazette*, November 8, 2003.
- [3] D. Huber and M. Hebert, "A New Approach to 3D Terrain Mapping", *IEEE/RSJ International Conference on Intelligent Robotics and Systems*, 1999.
- [4] D. Huber, "Automatic Three-dimensional Modeling from Reality", *Doctoral Dissertation, Carnegie Mellon University*, 2002.
- [5] D. Huber and M. Hebert, "3D Modeling Using a Statistical Sensor Model and Stochastic Search", *IEEE International Conference on Computer Vision and Pattern Recognition*, 2003.
- [6] D. Langer, M. Mettenleiter, F. Hartl and C. Frohlich, "Imaging Ladar for 3-D Surveying and CAD Modeling of Real World Environments", *International Journal of Robotics Research*, vol 19, no 11.
- [7] R. Madhavan, M. Dissanayake and H. Durrant-Whyte, "Autonomous underground navigation of an LHD using a combined ICP-EKF approach", *International Conference on Robotics and Automation*, 1998.
- [8] R. Madhavan, G. Dissanayake and H. Durrant-Whyte, "Map-building and map-based localization in an underground-mine by statistical pattern matching", *International Conference on Pattern Recognition*, 1998.
- [9] A. Morris, D. Kurth, W. Whittaker and Scott Thayer, "Case Studies of a Borehole Deployable Robot for Limestone Mine Profiling and Mapping", *International Conference on Field and Service Robotics*, 2003.
- [10] J.M. Roberts, E.S. Duff, P. Corke, P. Sikka, G.J. Winstanley and J. Cunningham, "Autonomous Control of Underground Mining Vehicles using Reactive Navigation", *International Conference on Robotics and Automation*, 2000.
- [11] S. Scheduling, E. Nebot, M. Stevens, H. Durrant-Whyte, J. Roberts, P. Corke, J. Cunningham, and B. Cook, "Experiments in autonomous underground guidance", *IEEE International Conference on Robotics and Automation*, 1997.
- [12] S. Scheduling, G. Dissanayake, E.M. Nebot, and H. Durrant-Whyte, "An Experiment in Autonomous Navigation of an Underground Mining Vehicle", *IEEE Transactions on Robotics and Automation*, vol 15, no 1, 1999.
- [13] G. Shaffer, A. Stentz, W. Whittaker, and K. Fitzpatrick, "Position estimator for underground mine equipment", *IEEE Transactions on Industry Applications*, Volume: 28 Issue: 5, Sep/Oct 1992.
- [14] S. Thrun et al., "A System for Volumetric Robotic Mapping of Abandoned Mines", *International Conference on Robotics and Automation*, 2003.